Discriminating between similar languages in Twitter using label propagation

Will Radford
Xerox Research Centre Europe
6 chemin de Maupertuis
38240 Meylan, France
will.radford@xrce.xerox.com

Matthias Gallé
Xerox Research Centre Europe
6 chemin de Maupertuis
38240 Meylan, France
matthias.galle@xrce.xerox.com

arXiv:1607.05408v1 [cs.CL] 19 Jul 2016

ABSTRACT

Identifying the language of social media messages is an important first step in linguistic processing. Existing models for Twitter focus on content analysis, which is successful for dissimilar language pairs. We propose a label propagation approach that takes the social graph of tweet authors into account as well as content to better tease apart similar languages. This results in state-of-the-art shared task performance of 76.63%, 1.4% higher than the top system.

Categories and Subject Descriptors

I.2.7 [**Artificial Intelligence**]: Natural Language Processing

Keywords

Language identification, Social media, Label propagation

1. INTRODUCTION

Language identification is a crucial first step in textual data processing and is considered feasible over formal texts [4]. The task is harder for social media (e.g. Twitter) where text is less formal, noisier and can be written in wide range of languages. We focus on identifying *similar* languages, where surface-level content alone may not be sufficient. Our approach combines a content model with evidence propagated over the social network of the authors. For example, a user well-connected to users posting in a language is more likely to post in that language. Our system scores 76.63%, 1.4% higher than the top submission to the `tweetLID` workshop.[1]

2. BACKGROUND

Traditional language identification compares a document with a language fingerprint built from n-gram bag-of-words (character or word level). Tweets carry additional metadata useful for identifying language, such as geolocation [3], username [2, 3] and URLs mentioned in the tweet [2].

Other methods expand beyond the tweet itself to use a histogram of previously predicted languages, those of users `@`-mentioned and lexical content of other tweets in a discussion [3]. Discriminating between similar languages was the focus of the VarDial workshop [7], and most submissions used content analysis. These methods make limited use of the social context in which the authors are tweeting – our research question is "Can we identify the language of a tweet using the social graph of the tweeter?".

Label propagation approaches [8] are powerful techniques for semi-supervised learning where the domain can naturally be described using an undirected graph. Each node contains a probability distribution over labels, which may be empty for unlabelled nodes, and these labels are propagated over the graph in an iterative fashion. Modified Adsorption (MAD)[6], is an extension that allows more control of the random walk through the graph. Applications of LP and MAD are varied, including video recommendation [1] and sentiment analysis over Twitter [5].

3. METHOD

Our method predicts the language ℓ for a tweet t by combining scores from a content model and a graph model that takes social context into account, as per Equation 1:

$$lang(t) = \operatorname*{argmax}_{\ell} \lambda_1 p(\ell|t, \theta_{content}) + \lambda_2 p(\ell|t, \theta_{social}) \quad (1)$$

Where $\theta_{content}$ are the content model parameters, θ_{social} the social model parameters.[2]

3.1 Content model

Our content model is a 1 vs. all ℓ_2 regularised logistic regression model[3] with character 2- to 5-grams features, not spanning over word boundaries. The scores for a tweet are normalised to obtain a probability distribution.

3.2 Social model

We use a graph to model the *social media* context, relating tweets to one another, authors to tweets and other authors. Figure 1 shows the graph, composed of three types of nodes: tweets (T), users (U) and the "world" (W). Edges are created between nodes and weighted as follows: **T-T** the unigram cosine similarity between tweets, **T-U** weighted 100 between

[1] `http://komunitatea.elhuyar.org/tweetlid`

[2] We do not optimise λ_1 and λ_2, setting them to 0.5.

[3] We use `scikit-learn`: http://scikit-learn.org

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...$15.00.

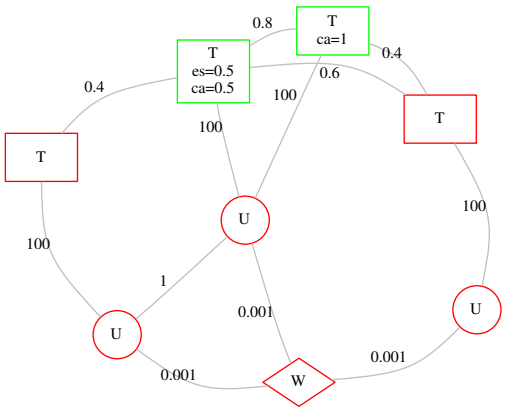

Figure 1: Graph topology. Rectangular nodes are tweets, circular nodes are users and the diamond represents the world. Some tweet nodes are labelled with an initial distribution over language labels and others are unlabelled.

a tweet and its author, **U-U** weighted 1 between two users in a "follows" relationship and **U-W** weighted 0.001 to ensure a connected graph for the MAD algorithm.

We create the graph using all data, and training set tweets have an initial language label distribution.[4] A naïve approach to building the tweet-tweet subgraph requires $O(n^2)$ comparisons, measuring the similarity of each tweet with all others. Instead, we performed k-nearest-neighbour classification on all tweets, represented as a bag of unigrams, and compared each tweet and the top-k neighbours.[5] We use Junto (MAD) [6] to propagate labels from labelled to unlabelled nodes. Upon convergence, we renormalise label scores for initially unlabelled nodes to find the value of θ_{graph}.

4. EVALUATION

The `tweetLID` workshop shared task requires systems to identify the language of tweets written in Spanish (ES), Portuguese (PT), Catalan (CA), English (EN), Galician (GL) and Basque (EU). Some language pairs are similar (ES and CA; PT and GL) and this poses a challenge to systems that rely on content features alone. We use the supplied evaluation corpus, which has been manually labelled with six languages and evenly split into training and test collections. We use the official evaluation script and report precision, recall and F-score, macro-averaged across languages. This handles ambiguous tweets by permitting systems to return any of the annotated languages. Table 1 shows that using the content model alone is more effective for languages that are distinct in our set of languages (i.e. English and Basque). For similar languages, adding the social model helps discriminate them (i.e. Spanish, Portuguese, Catalan and Galician), particularly those where a less-resourced language is similar to a more popular one. Using the social graph almost doubles the F-score for undecided (UND) languages, either not in the set above or hard-to-identify, from 18.85% to 34.95%. Macro-averaged, our system scores 76.63%, higher than the best score in the competition: 75.2%.

[4]We assume a uniform distribution for AMB tweets.
[5]We used `scikit-learn` with $k = 0.25 * ntweets$.

	Content			Content + Social		
	P	R	F	P	R	F
ES ♢	92.64	95.69	94.14	93.55	95.89	**94.70**
PT ♠	89.81	92.58	91.17	94.87	92.52	**93.68**
CA ♢	81.14	87.19	84.06	85.22	90.17	**87.62**
EN	77.42	76.18	**76.79**	77.86	70.53	74.01
GL ♠	56.93	52.93	54.85	65.15	50.35	**56.80**
EU	92.41	76.29	**83.58**	94.41	68.01	79.06
amb	100.00	89.56	**94.49**	100.00	85.54	92.21
und	66.67	10.98	18.85	45.06	28.54	**34.95**
avg	82.13	72.67	74.74	82.01	72.69	**76.63**

Table 1: Experimental results. ♢/♠ are similar pairs.

5. CONCLUSION

Our approach uses social information to help identify the language of tweets. This shows state-of-the-art performance, especially when discriminating between similar languages. A by-product of our approach is that users are assigned a language distribution, which may be useful for other tasks.

6. REFERENCES

[1] S. Baluja, R. Seth, D. Sivakumar, Y. Jing, J. Yagnik, S. Kumar, D. Ravichandran, and M. Aly. Video suggestion and discovery for youtube: Taking random walks through the view graph. In *Procs. of the 17th International Conference on World Wide Web*, WWW '08, pages 895–904, New York, NY, USA, 2008. ACM.

[2] S. Bergsma, P. McNamee, M. Bagdouri, C. Fink, and T. Wilson. Language identification for creating language-specific twitter collections. In *Procs. of the Second Workshop on Language in Social Media*, LSM '12, pages 65–74, Stroudsburg, PA, USA, 2012. ACL.

[3] S. Carter, W. Weerkamp, and M. Tsagkias. Microblog language identification: Overcoming the limitations of short, unedited and idiomatic text. *Lang. Resour. Eval.*, 47(1):195–215, Mar. 2013.

[4] P. McNamee. Language identification: A solved problem suitable for undergraduate instruction. *J. Comput. Sci. Coll.*, 20(3):94–101, Feb. 2005.

[5] M. Speriosu, N. Sudan, S. Upadhyay, and J. Baldridge. Twitter polarity classification with label propagation over lexical links and the follower graph. In *Procs. of the First workshop on Unsupervised Learning in NLP*, pages 53–63, Edinburgh, Scotland, July 2011. ACL.

[6] P. P. Talukdar and K. Crammer. New regularized algorithms for transductive learning. In *Procs. of the European Conference on Machine Learning and Knowledge Discovery in Databases: Part II*, ECML PKDD '09, pages 442–457, Berlin, Heidelberg, 2009. Springer-Verlag.

[7] M. Zampieri, L. Tan, N. Ljubešić, and J. Tiedemann. A report on the dsl shared task 2014. In *Proceedings of the First Workshop on Applying NLP Tools to Similar Languages, Varieties and Dialects*, pages 58–67, Dublin, Ireland, August 2014. ACL and DCU.

[8] X. Zhu and Z. Ghahramani. Learning from labeled and unlabeled data with label propagation. Technical Report CMU-CALD-02-107, CMU, 2002.